

AD-776 610

THE VALUE OF SEQUENTIAL INFORMATION

Allen Clinton Miller, III

Stanford University

Prepared for:

Office of Naval Research
Advanced Research Projects Agency
National Science Foundation

31 January 1974

DISTRIBUTED BY:

NTIS

National Technical Information Service
U. S. DEPARTMENT OF COMMERCE
5285 Port Royal Road, Springfield Va. 22151


AD776610

Research Report No. EES-DA-73-2
January 1974

THE VALUE OF SEQUENTIAL INFORMATION

ALLEN C. MILLER III

DECISION ANALYSIS PROGRAM
Professor Ronald A. Howard
Principal Investigator



DEPARTMENT OF ENGINEERING-ECONOMIC SYSTEMS
Stanford University
Stanford, California 94305

Reproduced by
NATIONAL TECHNICAL
INFORMATION SERVICE
U S Department of Commerce
Springfield VA 22151

SPONSORSHIPS

Advanced Research Projects Agency, Human Resources Research Office
ARPA Order No. 2449, monitored by Engineering Psychology Programs,
Office of Naval Research, under Contract No. N00014-67-A-0112-0077
(NR 197-024) covering the period May 1, 1973 to December 30, 1973.

National Science Foundation under NSF Grant GK-36491.

Approved for public release with distribution unlimited; reproduction in whole, or in part, permitted for any purpose of the United States Government.



REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER EES-DA-73-2	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) "The Value of Sequential Information"		5. TYPE OF REPORT & PERIOD COVERED Technical 5/1/73 to 12/30/73
7. AUTHOR(s) Allen Clinton Miller, III		6. PERFORMING ORG. REPORT NUMBER
9. PERFORMING ORGANIZATION NAME AND ADDRESS The Board of Trustees of the Leland Stanford Junior University, c/o Office of Research Admini- strator, Encina Hall, Stanford, California 94305		8. CONTRACT OR GRANT NUMBER(s) N00014-67-A-0112-0077
11. CONTROLLING OFFICE NAME AND ADDRESS Advanced Research Projects Agency Human Resources Research Office 1400 Wilson Blvd., Arlington, Virginia 22209		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS 000000 ARPA Order #2449
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office) Engineering Psychology Programs, Code 455 Office of Naval Research 800 N. Quincy Street, Arlington, VA 22217		12. REPORT DATE January 31, 1974
		13. NUMBER OF PAGES 24 pages
		15. SECURITY CLASS. (of this report) Unclassified
16. DISTRIBUTION STATEMENT (of this Report) Approved for public release; distribution unlimited.		18. DECLASSIFICATION/DOWNGRADING SCHEDULE
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)		
18. SUPPLEMENTARY NOTES A paper submitted to MANAGEMENT SCIENCE/THEORY.		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number) DECISION ANALYSIS DECISION THEORY SEQUENTIAL INFORMATION DECISION-MAKING INFORMATION VALUE OF INFORMATION		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) In decision analysis we normally consider the value of information to be a constant against which the cost of information is compared. However, when it is possible to buy information sequentially, the value of information is not a constant. Rather, it is a function of the prices of the various pieces of information, or "observables." When we are faced with a decision and learn one observable, this information not only helps us make the		

Block 20 (continued)

original decision, but also helps us decide if we should pay for more observables. For this reason, the first observable has a value above and beyond that which we would assign if there were no possibility of obtaining additional information. To decide whether or not to buy one observable we must know the prices of all the observables.

UNCLASSIFIED

THE VALUE OF SEQUENTIAL INFORMATION*

Allen C. Miller, III

Abstract

In decision analysis we normally consider the value of information to be a constant against which the cost of information is compared. However, when it is possible to buy information sequentially, the value of information is not a constant. Rather, it is a function of the prices of the various pieces of information, or "observables." When we are faced with a decision and learn one observable, this information not only helps us make the original decision, but also helps us decide if we should pay for more observables. For this reason, the first observable has a value above and beyond that which we would assign if there were no possibility of obtaining additional information. To decide whether or not to buy one observable we must know the prices of all the observables.

Introduction

Using decision analysis [1,4,6,7]** it is possible to calculate the value of one or more pieces of information--called "observables"--when a decision must be made in the face of uncertainty. This information has value because it can affect the decision and lead to a greater expected

*This research was partially supported by the Advanced Research Projects Agency of the Department of Defense, as monitored by the Office of Naval Research under Contract No. N00014-67-A-0112-0077, and by the National Science Foundation under NSF Grant GK-36491.

**Numbers in square brackets refer to the References.

profit. However the possibility of buying information sequentially presents the decision maker with a set of secondary decisions: which observables should he buy and in which order should he buy them? It is possible that knowing one observable affects not only the primary decision, but also the decision to buy additional information. In that case the value of knowing the first observable is greater than it would be if it affected only the primary decision. The prices of the observables affect the decision maker's willingness to buy additional information. For this reason the amount that the value of learning each observable is increased by the possibility of buying additional information depends on the prices of all the observables.

Notation

The notation used in this paper is an extension of Howard's inferential notation [2,3], which shows the decision maker's state of information explicitly. Although inferential notation is very useful, the following operator notation is somewhat clearer and more compact when dealing with the value of information.

A = a state of information on which probability assignments are made,

δ = the decision maker's prior state of information,

$\{x|A\}$ = the density function (for continuous random variables) or mass function (for discrete random variables) of the random variable x , given the state of information A ,

$\langle x|A \rangle = \int x \{x|A\} dx$ = the expected value of x given the state of information A ,

$\bar{x} = \langle x | \delta \rangle$ = the expected value of x given the decision maker's prior state of information ,

$\pi(x_1, \dots, x_n, c)$ = a profit function that depends on the state and control variables,

c = a control or decision variable (possibly vector-valued) upon which the profit depends ,

x_i = a state variable; a random variable upon which the profit depends.
(If we are dealing with perfect information, x_i is also an observable.)

y_j = an observable; a random variable whose actual value we have an opportunity to learn ,

$E_{x_i} \pi(x_1, \dots, x_n, c) = \int \pi(x_1, \dots, x_n, c) \{x_i | \delta\} dx$ = the expected value of the profit with respect to x_i ,

$\max_c \pi(x_1, \dots, x_n, c)$ = the maximum of the profit with respect to c ,

K_{x_i} = cost of learning x_i ,

$V_{x_i}^N$ = the value of learning x_i given that no additional information is purchased; the value of individual information about x_i (the superscript N means "no additional information") ,

$V_{x_i x_j \dots x_r}^N$ = the value of learning x_i, x_j, \dots , and x_r simultaneously given that no additional information is purchased; the value of simultaneous information about x_i, x_j, \dots , and x_r ,

$V_{x_i}^R = V_{x_1 \dots x_n}^N - K_{x_1} - \dots - K_{x_{i-1}} - K_{x_{i+1}} - \dots - K_{x_n}$, the residual value of learning x_i when all of the available information is purchased simultaneously; the residual value of x_i , (the superscript R stands for "residual value") ,

V_{x_i} = the value of learning x_i when additional information can be purchased; the value of sequential information about x_i .

Strings of expectation and maximization operators are interpreted as follows:

$$\max_c E_x \pi(x, c) = \max_c \left(\int_{-\infty}^{\infty} \pi(x, c) \{x|\delta\} dx \right)$$

$$D_y \max_c E_x \pi(x, c) = \max_c \left(\int_{-\infty}^{\infty} \pi(x, c) \{x|y, \delta\} dx \right)$$

$$E_y \max_c E_x \pi(x, c) = \int_{-\infty}^{\infty} \max_c \left(\int_{-\infty}^{\infty} \pi(x, c) \{x|y, \delta\} dx \right) \{y|\delta\} dy$$

D_y is a modifier that conditions the expectations in its argument on y . E_x is interpreted as the expected value with respect to x given all the control or random variables that appear as subscripts of operators or modifiers that include E_x in their argument.

In this paper we assume that none of the probability density functions depends on the control variable.

Types of Sequential Information Problems

A decision problem involving sequential information can be placed in one of several categories depending on whether or not the problem has each of the following properties.

1. Additive prices. The prices of the observables are additive

when the cost of learning any group of observables simultaneously equals the sum of the costs of learning each observable separately.

2. Certain prices. If the decision maker knows the price of each observable and group of observables, the problem can be analyzed in terms of certain prices.
3. Perfect information. When a piece of information reveals the exact value of one or more of the arguments of the profit function, it is called perfect information.

For simplicity, we do not consider utility and risk preference in this paper. However the extension of the following results to include these concepts is straightforward.

The following example introduces ideas of sequential information with the simplest type of problem: the case of additive, certain prices for perfect information. After the example the results are generalized to apply to all types of sequential information problems.

An Example: A Bidding Problem Revisited

The example addresses the problem of submitting a single, sealed bid for a contract when faced with uncertainty about the lowest competing bid and the production cost required to fulfill the contract. This problem is discussed in some detail by Howard [2,3].

Suppose our company's objective in bidding is to maximize its expected profit. Let p be our production cost for performing the contract, let l be our competitors' lowest bid, and let b be our company's bid. If our bid exceeds the lowest competing bid, our company

will not win the contract and our net profit will be zero. However if we submit the lowest bid, our company will win the contract and make a profit equal to the difference between our bid and the production cost. Thus our company's profit, π , is defined by

$$\pi(p, \ell, b) = \begin{cases} b - p & : b < \ell \\ 0 & : b \geq \ell \end{cases}$$

We do not know the exact values of p and ℓ , but we can assign prior probability density functions -- $\{p|\xi\}$ and $\{\ell|\xi\}$ -- based on our company's experience with similar bidding situations. To simplify the computations required for this example we shall assume that $\{p|\xi\}$ and $\{\ell|\xi\}$ have been assessed as independent, uniform density functions, with p ranging from zero to one and ℓ ranging from zero to two. These distributions are shown in Fig. 1.

Consider the possibility of buying perfect information about p and ℓ at costs of K_p and K_ℓ respectively. Howard [2] has shown that the values of individual and simultaneous information about p and p and ℓ are :

$$V_p^N = (E \max_p E - \max_b E E) \pi(p, \ell, b) = 1/96$$

$$V_\ell^N = (E \max_\ell E - \max_b E E) \pi(p, \ell, b) = 27/96$$

$$V_{p\ell}^N = (E E \max_{p\ell} - \max_b E E) \pi(p, \ell, b) = 29/96$$

Howard points out that $V_{p\ell}^N$ does not equal the sum of V_p^N and V_ℓ^N . In other words, the value of learning both p and ℓ is greater than the sum of the values of learning p and ℓ individually.*

* It can be shown that, in general, the value of learning several observables simultaneously can be greater than or less than the sum of the values of learning each observable individually [5].




Figure 1. Probability density functions for production cost and lowest competing bid

By comparing the expected profits associated with learning p and l individually and simultaneously, we find the decision regions shown in Fig. 2. As Fig. 2 shows, it is necessary to know the prices of both pieces of information to decide which observable to buy, even when only individual and simultaneous purchases are possible. When we consider sequential information it will be possible to draw a diagram similar to Fig. 2, but the decision regions become more complicated.

To determine the value of sequential information about p , assume for a moment that we already know p and are trying to decide whether or not to pay to learn l . If we decide not to learn l , then our maximum expected profit is:

$$\max_b E_l \pi(p, l, b) = (1 - p/2)^2/2$$

On the other hand, if we decide to learn l , our maximum expected profit is

$$E_l \max_b \pi(p, l, b) = (1 - p/2)^2$$

The increase in expected profit caused by learning l after we know p is

$$(E_l \max_b - \max_b E_l) \pi(p, l, b) = (1 - p/2)^2/2$$

We are willing to pay for l when the increase in expected profit exceeds K_l . For certain values of K_l , the decision to buy perfect information about l depends on the value of p that we learned earlier, which means that learning p can help us decide whether to buy more information as well as how to bid. The maximum expected profit when we know that will receive perfect information about p with an option to pay K_l for l is




Figure 2. Decision regions for individual and simultaneous information

$$\begin{aligned}
 & E_p \max \left\{ \begin{array}{l} E_l \max_b \pi(p, l, b) - K_l \\ \max_b E_l \pi(p, l, b) \end{array} \right\} \\
 &= \left\{ \begin{array}{l} E_p (1 - p/2)^2 - K_l : K_l \leq 1/8 \\ 2(1-\sqrt{2K_l}) \int_0^{1-\sqrt{2K_l}} [(1-p/2)^2 - K_l] \{p|\delta\} dp \\ \quad + \int_{1-\sqrt{2K_l}}^1 [(1-p/2)^2/2] \{p|\delta\} dp : 1/8 < K_l < 1/2 \\ E_p (1 - p/2)^2/2 : 1/2 \leq K_l \end{array} \right\} \\
 &= \left\{ \begin{array}{l} 56/96 - K_l : K_l \leq 1/8 \\ 60/96 + (4/3)K_l \sqrt{2K_l} - 2K_l : 1/8 < K_l < 1/2 \\ 28/96 : 1/2 \leq K_l \end{array} \right\}
 \end{aligned}$$

The corresponding increase in expected profit is the value of sequential information about p .

$$\begin{aligned}
 V_p(K_l) &= E_p \max \left\{ \begin{array}{l} E_l \max_b \pi(p, l, b) - K_l \\ \max_b E_l \pi(p, l, b) \end{array} \right\} - \max_b E_p E_l \pi(p, l, b) \\
 &= \left\{ \begin{array}{l} 29/96 - K_l : K_l \leq 1/8 \\ 33/96 + (4/3)K_l \sqrt{2K_l} - 2K_l : 1/8 < K_l < 1/2 \\ 1/96 : 1/2 \leq K_l \end{array} \right\}
 \end{aligned}$$

By carrying out a similar calculation we can find the maximum expected profit when we know that we will receive perfect information about ℓ with an option to pay K_p for p . This expected profit is

$$\begin{aligned}
 & E_{\ell} \max \left\{ \begin{array}{l} E_p \max_b \pi(p, \ell, b) - K_p \\ \max_b E \pi(p, \ell, b) \end{array} \right\} \\
 &= \left\{ \begin{array}{l} \int_{\sqrt{2K_p}}^{1-\sqrt{2K_p}} (\ell^2/2 - K_p) \{ \ell | \delta \} d\ell + \int_{1-\sqrt{2K_p}}^2 (\ell - 1/2) \{ \ell | \delta \} d\ell : K_p < 1/8 \\ \int_{1/2}^2 (\ell - 1/2) \{ \ell | \delta \} d\ell : K_p \geq 1/8 \end{array} \right\} \\
 &= \left\{ \begin{array}{l} 56/96 + (2/3)K_p\sqrt{2K_p} - K_p/2 : K_p < 1/8 \\ 54/96 : K_p \geq 1/8 \end{array} \right\}
 \end{aligned}$$

The corresponding increase in expected profit is the value of sequential information about ℓ .

$$\begin{aligned}
 V_{\ell}(K_p) &= E_{\ell} \max \left\{ \begin{array}{l} E_p \max_b \pi(p, \ell, b) - K_p \\ \max_b E \pi(p, \ell, b) \end{array} \right\} - \max_b E E \pi(p, \ell, b) \\
 &= \left\{ \begin{array}{l} 29/96 + (2/3)K_p\sqrt{2K_p} - K_p/2 : K_p < 1/8 \\ 27/96 : K_p \geq 1/8 \end{array} \right\}
 \end{aligned}$$

By comparing the expected profits associated with all the possible ways we can learn p and ℓ --individually, simultaneously and sequentially--we find the decision regions shown in Fig. 3. This figure also shows $V_p(K_{\ell})$ and $V_{\ell}(K_p)$ plotted against the appropriate axes, since




Figure 3. Decision regions for sequential information

these functions form boundaries of the decision regions. In this problem the option of buying information simultaneously is completely dominated by the other alternatives, even though there are pairs of prices such that buying both pieces of information simultaneously is preferable to buying either piece of information individually. A comparison of Figs. 2 and 3 shows that when information is available sequentially, we are willing to buy information at pairs of prices that were not advantageous when only individual and simultaneous information was available.

Furthermore, there are pairs of prices such that our best decision when offered individual and simultaneous information is to pay K_ℓ to learn ℓ , while our best initial decision when offered all types of information is to pay K_p to learn p .

General Properties of Sequential Information Problems with Additive, Certain Prices

When all observable prices are additive and certain, we can formulate a general sequential-information problem in terms of a set of state variables (x_1, \dots, x_m) and a set of observables (y_1, \dots, y_n) with the corresponding set of observable prices $(K_{y_1}, \dots, K_{y_n})$. When an observable is equal to one of the state variables it represents perfect information. However, by treating observables and state variables separately, we can also deal with imperfect information. Since the possibility of buying observables sequentially increases the value of both imperfect and perfect information, the value of imperfect information about a state variable can exceed the value of perfect, individual information about the same state variable.

As the preceding example demonstrates, the value of sequential information about an observable is, in general, a function of the prices of all of the other observables. Thus,

$$V_{y_i} = V_{y_i}(K_{y_1}, \dots, K_{y_{i-1}}, K_{y_{i+1}}, \dots, K_{y_n})$$

The dependence of V_{y_i} on the prices of the observables can be seen in the decision tree in Fig. 4. This tree shows all the information-purchasing decisions as well as the primary decision to choose a value for the control variable. The value of sequential information about y_i is equal to the difference between the expected profit associated with branch B_i (the branch where we first learn y_i) and the expected profit associated with branch B_0 (the branch where we do not learn any information. Using the tree in Fig. 4 we can find V_{y_i} algebraically.

$$V_{y_i} = E_{y_i} \max \left\{ \begin{array}{l} \max_c E_{x_1} \dots E_{x_m} \pi(x_1, \dots, x_m, c) \\ \max_{j \neq i} \left(E_{y_j} \max \left\{ \begin{array}{l} \max_c E_{x_1} \dots E_{x_m} \pi(x_1, \dots, x_m, c) \\ \max_{k \neq i, j} \left(E_{y_k} \max \{ \dots \} - K_{y_k} \right) \end{array} \right\} - K_{y_j} \right) \end{array} \right\} \\ - \max_c E_{x_1} \dots E_{x_m} \pi(x_1, \dots, x_m, c)$$

This function has several properties which were demonstrated in the preceding example. Specifically we can show that as any price increases, V_{y_i} must decrease or remain constant. However V_{y_i} cannot decrease by more than the price increase. Thus,

$$\frac{\partial V_{y_i}}{\partial K_{y_j}} \in [-1, 0]$$

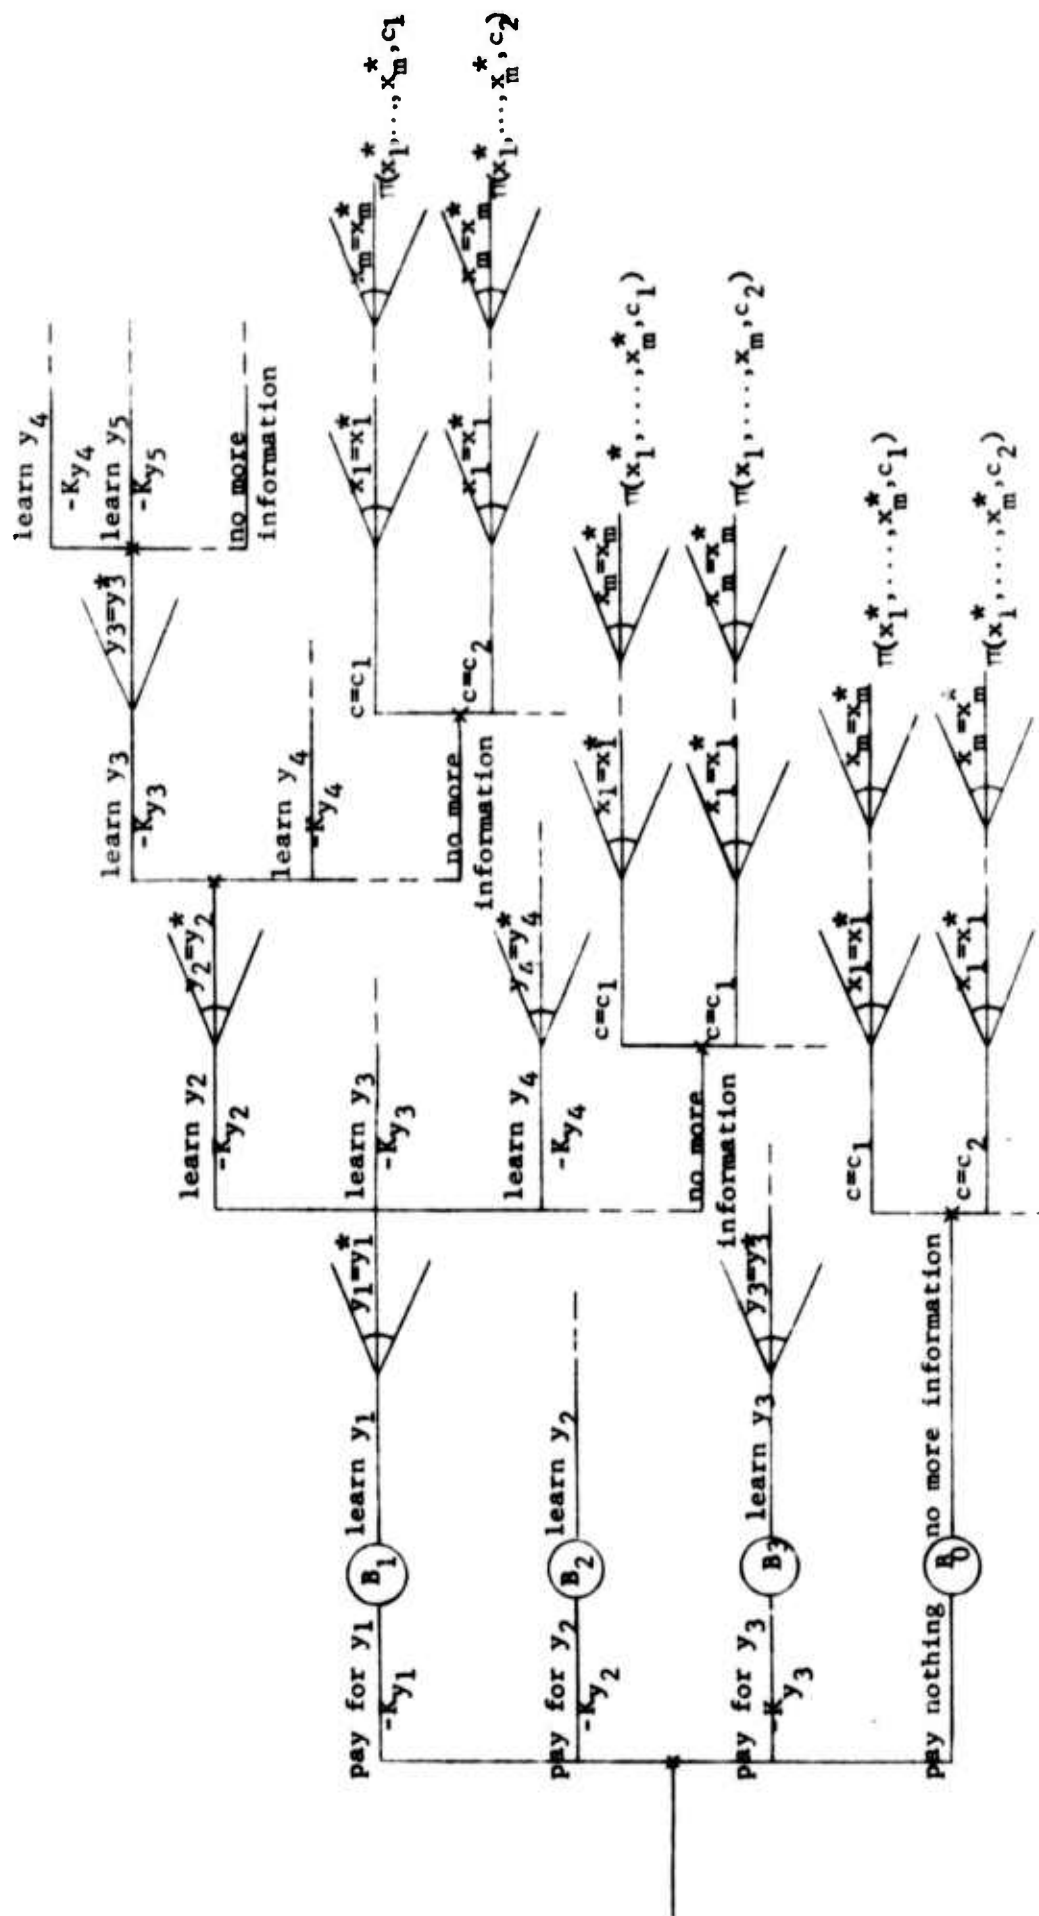


Figure 4. Sequential information decision tree for an expected profit decision maker

The proof of this statement [5] depends on the fact that all of the prices in the expression for V_{y_i} are subtracted from some quantity. The derivative of V_{y_i} becomes a series of expectations of quantities that are either minus one, zero, or some number between these limits.

It can also be shown that

$$\partial^2 V_{y_i} / \partial K_{y_j}^2 \geq 0$$

The proof [5] is based on the fact that, as K_{y_j} increases, it becomes less desirable for the decision maker to pay for the j^{th} observable regardless of the information he has already learned. In the expression for V_{y_i} this means that there are fewer values of previously-learned observables for which the expectations include a term with $-K_{y_j}$. Thus $\partial V_{y_i} / \partial K_{y_j}$ becomes an expectation of more terms that are zero and fewer terms that are minus one, so the first derivative must increase or remain constant when K_{y_j} increases.

The significance of these properties of the value of sequential information is that they allow V_{y_i} to exceed the corresponding values of individual and simultaneous information for certain sets of observable prices. It is easy to see that V_{y_i} must be at least as great as $V_{y_i}^N$ and $V_{y_i}^R$ for any set of observable prices.* When we buy information sequentially, we have the opportunity to buy just one observable or all of the observables. Thus we can always achieve the expected profits associated with individual and simultaneous information by making the

* Using $V_{y_i}^R(K_{y_1}, \dots, K_{y_{i-1}}, K_{y_{i+1}}, \dots, K_{y_n})$, we can describe the price pairs where we will buy simultaneous information in terms of a function similar to V_{y_i} . We are willing to buy all the observables simultaneously when $K_{y_i} < V_{y_i}^R$.

proper set of sequential decisions.

The bidding example showed that, for certain sets of observable prices, V_{y_i} can exceed both $V_{y_i}^N$ and $V_{y_i}^R$. This phenomenon, which makes it possible to buy sequential information at sets of prices that are not advantageous for individual or sequential purchases, occurs whenever learning y_i can affect our decision to learn other observables. Suppose that for some set of observable prices our decision to buy additional information changes when we learn different values of y_i . Therefore, for some value of y_i , our best decision must be to buy additional information. For this value of y_i , our expected profit must be greater when we pay for some other observable, y_j , than when we refuse additional information. Therefore,

$$D_{y_i} E_{y_j} \max \left\{ \begin{array}{l} \max_c E_{x_1} \dots E_{x_m} \pi \\ \max_{k \neq i, j} (E_{y_k} \max \{ \dots \} - K_{y_k}) \end{array} \right\} - K_{y_j} > D_{y_i} \max_c E_{x_1} \dots E_{x_m} \pi$$

Expanding this inequality and taking the expected value of both sides yields $V_{y_i} > V_{y_i}^N$. The proof that V_{y_i} exceeds $V_{y_i}^R$ is more complicated but essentially similar [5].

Observables with Uncertain Prices

If the prices of the observables are uncertain, we can assign prior probability density functions for each price. In general, the resulting distributions need not be independent, so our state of information is represented by a joint distribution on the state variables, observables,

and prices

$$\{x_1, \dots, x_m, y_1, \dots, y_n, K_{y_1}, \dots, K_{y_n} | \mathcal{E}\}$$

When we commit ourselves to paying for the i^{th} observable, we learn both y_i and K_{y_i} . After paying K_{y_i} we must then decide whether to buy additional information.

The expected value of learning y_i sequentially, when we are trying to maximize expected profit, and when the cost of information is uncertain, is

$$\begin{aligned} E_{y_i} E_{K_{y_i}} \max & \left\{ \begin{aligned} & \max_c E_{x_1} \dots E_{x_m} \pi(x_1, \dots, x_m, c) \\ & \max_{j \neq i} \left(E_{y_j} E_{K_{y_j}} \left[\max \left\{ \begin{aligned} & \max_c E_{x_1} \dots E_{x_m} \pi(x_1, \dots, x_m, c) \\ & \max_{k \neq i, j} (E_{y_k} E_{K_{y_k}} [\dots]) \end{aligned} \right\} - K_{y_j} \right] \right) \end{aligned} \right\} \\ & - \max_c E_{x_1} \dots E_{x_m} \pi(x_1, \dots, x_m, c) \end{aligned}$$

If K_{y_j} is independent of y_i and K_{y_i} , for all i and j , this expression reduces to $V_{y_i}^N(\bar{K}_{y_1}, \dots, \bar{K}_{y_{i-1}}, \bar{K}_{y_{i+1}}, \dots, \bar{K}_{y_n})$. In other words we can use the same decision rules for buying information that we would with certain prices, except we use the expected values of the prices instead of the prices themselves.

On the other hand, if K_{y_j} is dependent on one of the other prices or the corresponding observables, we cannot characterize each uncertain price by its expected value. The maximizations required to determine the value of sequential information depend on quantities like

$$D_{y_i} D_{K_{y_i}} E_{y_j} K_{y_j}, \text{ which in turn require a knowledge of the entire joint}$$

density function for the state variables, observables, and prices. Since dependencies between these variables cannot generally be represented by a finite set of numbers, it is impossible to represent the decision rules in a finite-dimensional Euclidean space. However if we know the joint density function, we can determine the optimum strategy for buying information sequentially for that particular density function by looking at all of the ways we could buy the observables.

Observables with Non-Additive Prices

If the observable prices are certain but not additive, it is necessary to define a large number of prices, one for each of the ways we could learn each observable. Some of the prices are:

- K_{y_1} = cost of learning y_1 when none of the other observables are known;
- $K_{y_1|y_j}$ = cost of learning y_1 when y_j is known;
- $K_{y_1|y_j y_k \dots y_r}$ = cost of learning y_1 when y_j, y_k, \dots are known;
- $K_{y_1 y_j \dots y_r}$ = cost of learning y_1, y_j, \dots simultaneously when none of the observables are known and no additional information will be purchased.

Since each of these prices represents one dimension of a diagram such as that in Fig. 3, it is clear that we cannot hope to visualize the decision rules for buying information. However we can still describe the decision rules algebraically.

The decision rules can be simplified if we restrict the prices of the observables such that they can be described with relatively few numbers. One such restriction is to assume that the i^{th} observable costs

K_{y_i} if it is the first piece of information that we learn, and $\lambda_{y_i} K_{y_i}$ otherwise. The reduction factor λ_{y_i} represents the savings that result from having previously set up an information-gathering process to learn a different observable. In this case we can visualize the decision rules in terms of a diagram, such as that in Fig. 3, with the boundary representing V_{y_i} stretched along the K_{y_i} axes by a factor of $(1/\lambda_{y_i})$.

Conclusions

We cannot regard the value of learning one observable by itself as the maximum that we would be willing to pay for that piece of information. When it is possible to buy additional information, the value of the first observable may increase. How much it increases depends on the prices of the other observables, so it is necessary to know the prices of all the observables before we can decide whether to buy one of them. The decision to buy a piece of information must take into account all of the ways we can purchase the information individually, simultaneously, and sequentially.

References

- [1] Howard, R.A., "Decision Analysis: Applied Decision Theory," Fourth International Conference on Operational Research, Boston, 1966.
- [2] Howard, R.A., "Information Value Theory," IEEE Transactions on Systems Science and Cybernetics, Vol. SSC-2, No. 1, August 1966.
- [3] Howard, R.A., "Value of Information Lotteries," IEEE Transactions on Systems Science and Cybernetics, Vol. SSC-3, June 1967.
- [4] IEEE Transactions of Systems Science and Cybernetics, Special Issue on Decision Analysis, Vol. SCC-4, No. 3, September 1968.
- [5] Miller, A.C., "The Value of Sequential Information," Ph.D. dissertation, Stanford University, 1973.
- [6] Raiffa, H., and R.O. Schlaifer, "Applied Statistical Decision Theory," Harvard Graduate School of Business Administration, Harvard University, 1961.
- [7] Raiffa, H., "Decision Analysis: Introductory Lectures on Choices under Uncertainty," Addison-Wesley, 1968.